Layering as Optimization Decomposition
And Advances in Network Utility Maximization

Mung Chiang
Electrical Engineering Department, Princeton University

Internet 0 Meeting, MIT Media Lab
October 1, 2004

Motivation

Is the entire protocol stack a distributed optimization machine?

Are there alternative protocol stacks we can design?

Overview

- NUM Framework
- Recent Advance 1: Inelastic Traffic
- Recent Advance 2: Coupled Utilities
- Recent Advance 3: Wireless Network Power Control
- Recent Advance 4: Physical Layer Channel Coding
- Layering as Optimization Decomposition

Thanks To: S. Boyd, R. Calderbank, J. Doyle, M. Fazel, P. Hande, J. W. Lee, Y. Li, S. H. Low, D. Palomar, S. Zhang

NUM Framework

Beyond (linear) Network Flow Problems. Basic version:

$$
\begin{array}{ll}
\text{maximize} & \sum_s U_s(x_s) \\
\text{subject to} & \sum_{s:l\in L(s)} x_s \le c_l, \ \ \forall l, \\
& \mathbf{x} \succeq 0
\end{array}
$$

NUM Framework

Beyond (linear) Network Flow Problems. Basic version:

$$
\begin{array}{ll}
\text{maximize} & \sum_s U_s(x_s) \\
\text{subject to} & \sum_{s:l \in L(s)} x_s \leq c_l, \;\; \forall l, \\
& \mathbf{x} \succeq 0
\end{array}
$$

Current applications:

- Reverse engineering: TCP congestion control and Internet rate allocation (Steven's talk)
- Forward engineering: Network resource allocation, *e.g.*, power control
- Both reverse and forward engineering: TCP/IP/MAC/PHY interactions (Steven's talk)

Major approaches:

- Optimization-theoretic: distributed optimal solution algorithm
- Game-theoretic: Nash equilibrium characterization

Goals and Steps

Layering as NUM decomposition for:

- Internet 0
- Existing networks: FAST Copper (with S. Fraser and J. Cioffi)

But many pieces are missing in basic version of NUM

- What is utility?
- What is a channel?

Many encouraging progress very recently made

Inelastic Traffic

Shenker (1995): "Because of this elasticity, they can decrease their transmission rate in the presence of congestion." What about inelastic flows?

Inelastic Traffic

Shenker (1995): "Because of this elasticity, they can decrease their transmission rate in the presence of congestion." What about inelastic flows?

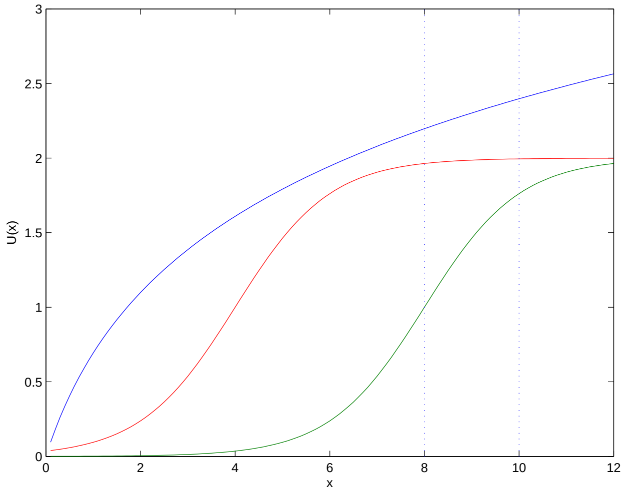

Inelastic Traffic

Shenker (1995): "Because of this elasticity, they can decrease their transmission rate in the presence of congestion." What about inelastic flows?

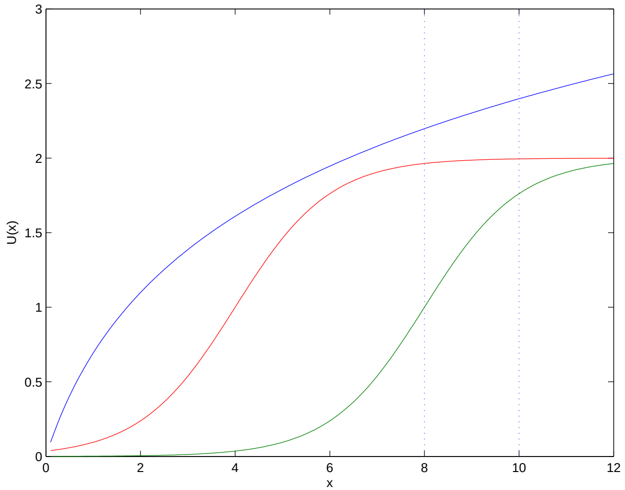

Rockafellar (1993): "The great watershed in optimization isn't between linearity and nonlinearity, but convexity and nonconvexity." Can we do nonconcave utility maximization?

When Does Canonical Distributed Algorithm Work?

- Set of bottleneck links does not change at optimal prices ($\Rightarrow$ continuity)
- $U^*(\mathbf{c})$ is a concave function ($\Rightarrow$ zero duality gap)

When Does Canonical Distributed Algorithm Work?

- Set of bottleneck links does not change at optimal prices ($\Rightarrow$ continuity)
- $U^*(\mathbf{c})$ is a concave function ($\Rightarrow$ zero duality gap)

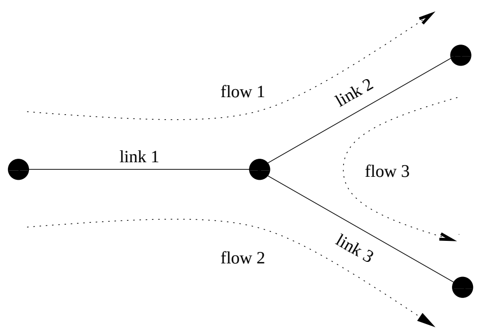

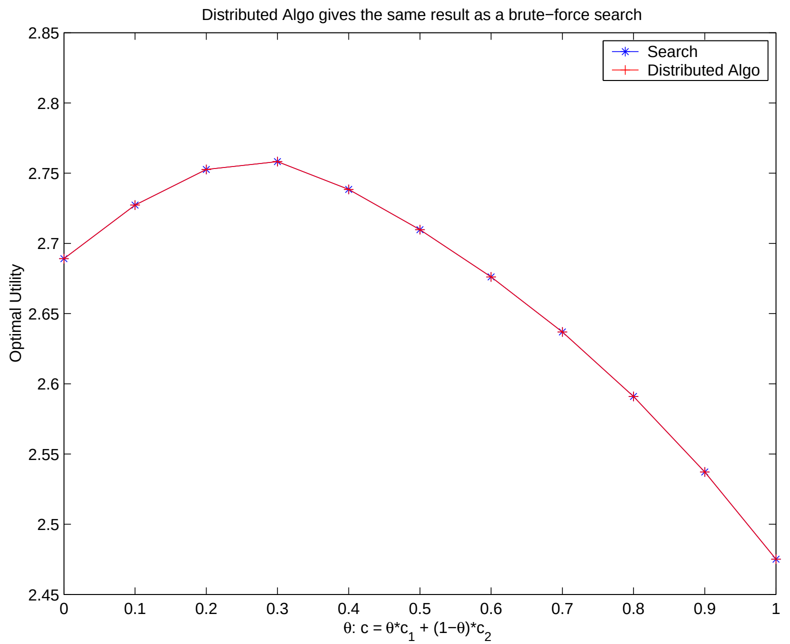

Inelastic Traffic: General Case

What can we do when duality gap is nonzero?

- Analysis: Can a certain amount of network utility be attained?
- Design: A new distributed algorithm?

Inelastic Traffic: General Case

What can we do when duality gap is nonzero?

- Analysis: Can a certain amount of network utility be attained?
- Design: A new distributed algorithm?

SOS Tool applied to inelastic flow congestion control: very numerically efficient (but centralized off-line) analysis

Obtain the globally optimal network utility with a second order polynomial link pricing

Inelastic Traffic: General Case

What can we do when duality gap is nonzero?

- Analysis: Can a certain amount of network utility be attained?
- Design: A new distributed algorithm?

SOS Tool applied to inelastic flow congestion control: very numerically efficient (but centralized off-line) analysis

Obtain the globally optimal network utility with a second order polynomial link pricing

A new distributed algorithm for link price update (by simple bisection) always converges to global optimality

Polynomial time in number of end hosts, exponential time in number of links (work well for small networks with many users)

Coupled Utilities: Formulations

U_s: a function of $\{x_i\}_{i \in I(s)}$ where $I(s)$ is the set of other sources whose rates is of concern to source s

- A hybrid model of selfish and non-selfish utilities
- Overlay or sensor network where some source nodes form a cluster
- DSL spectrum management and wireless power control

Coupled Utilities: Formulations

U_s: a function of $\{x_i\}_{i \in I(s)}$ where $I(s)$ is the set of other sources whose rates is of concern to source s

- A hybrid model of selfish and non-selfish utilities
- Overlay or sensor network where some source nodes form a cluster
- DSL spectrum management and wireless power control

$$
\begin{array}{ll}
\text{max} & U_1(x_1, x_2) + U_2(x_2, x_1, x_3) + U_3(x_3, x_1) \\
\text{s.t.} & x_1 + x_2 \leq c_1, \\
 & x_1 + x_3 \leq c_2, \\
 & \mathbf{x} \succeq 0
\end{array}
$$

Coupled Utilities: Formulations

U_s: a function of $\{x_i\}_{i \in I(s)}$ where $I(s)$ is the set of other sources whose rates is of concern to source s

- A hybrid model of selfish and non-selfish utilities
- Overlay or sensor network where some source nodes form a cluster
- DSL spectrum management and wireless power control

$$
\begin{array}{ll}
\text{max} & U_1(x_1, x_2) + U_2(x_2, x_1, x_3) + U_3(x_3, x_1) \\
\text{s.t.} & x_1 + x_2 \le c_1, \\
 & x_1 + x_3 \le c_2, \\
 & \mathbf{x} \succeq 0
\end{array}
$$

$$
\begin{array}{ll}
\text{max.} & U_1(x_1, y_1) + U_2(x_2, y_2, z_2) + U_3(x_3, y_3) \\
\text{s.t.} & y_1 = x_2, \\
 & y_2 = x_1, \\
 & z_2 = x_3, \\
 & y_3 = x_1, \\
 & x_1 + x_2 \le c_1, \\
 & x_1 + x_3 \le c_2, \\
 & \mathbf{x} \succeq 0
\end{array}
$$

Coupled Utilities: Optimal Distributed Algorithm

Associate extra constraints with consistency pricing: $\gamma_{12}, \gamma_{21}, \gamma_{23}, \gamma_{31}$

Decompose the Lagrangian into three sets, one for each source, *e.g.*:

$$U_1(x_1, y_1) + \gamma_{12} y_1 - \gamma_{21} x_1 - \gamma_{31} x_1 - \lambda_1 x_1 - \lambda_2 x_1$$

Maximizing over x_1 and y_1 locally at source 1

Coupled Utilities: Optimal Distributed Algorithm

Associate extra constraints with consistency pricing: $\gamma_{12}, \gamma_{21}, \gamma_{23}, \gamma_{31}$

Decompose the Lagrangian into four sets, one for each source, *e.g.*:

$$U_1(x_1, y_1) + \gamma_{12} y_1 - \gamma_{21} x_1 - \gamma_{31} x_1 - \lambda_1 x_1 - \lambda_2 x_1$$

Maximizing over x_1 and y_1 locally at source 1

Link price $\boldsymbol{\lambda}$ implicitly feed back by queue management process, updated the same way as in canonical distributed algorithm

Consistency price $\boldsymbol{\gamma}$ communicated through the local communication channels among the sources whose utilities depend on each other's rates. Consistency prices update is easy, *e.g.*:

$$\gamma_{12}(t+1) = \gamma_{12}(t) + \frac{\beta}{t}(y_1(t) - x_2(t))$$

Wireless Network Power Control

Link capacity depends on transmit powers and channel conditions

Nonlinear dependency, *e.g.*, energy efficiency or SIR as a function of transmit power

Wireless Network Power Control

Link capacity depends on transmit powers and channel conditions

Nonlinear dependency, *e.g.*, energy efficiency or SIR as a function of transmit power

Joint control of end-to-end rate allocation and per-hop power control

Recent results obtained for various cases:

- Cellular downlink case
- Cellular general case
- Cellular end-to-end case
- Multihop case with voluntary relays
- Multihop case with relay charges
- Incremental estimation in sensor networks

Multihop Case

s: index for sources

l: index for links

$L(s)$: set of links used by s

x_s: rate from source s

$$\begin{array}{ll} \text{maximize} & \sum_s U_s(x_s) \\ \text{subject to} & \sum_{s:l\in L(s)} x_s \le \log(K\,\mathsf{SIR}_l(\mathbf{P})), \ \ \forall l, \\ & \mathbf{x}, \mathbf{P} \succeq 0 \end{array}$$

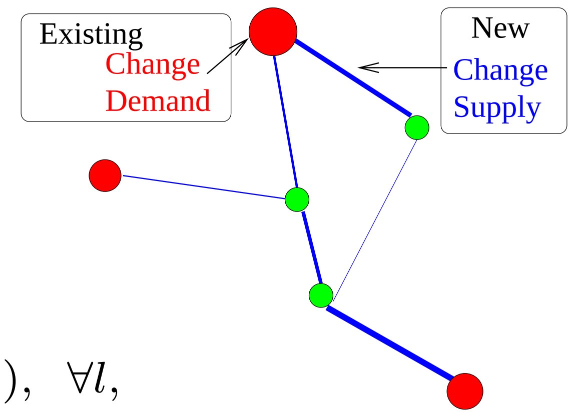

Multihop Case

s: index for sources
l: index for links
$L(s)$: set of links used by s
x_s: rate from source s

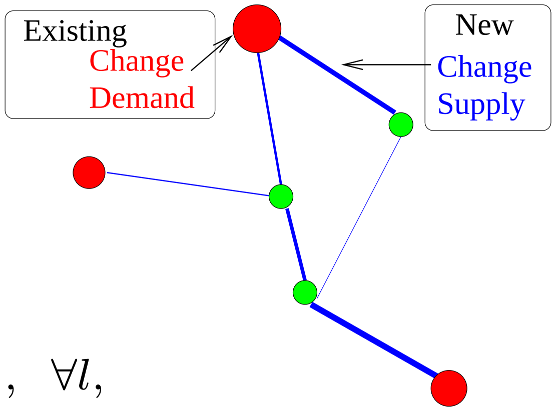

$$\begin{array}{ll} \text{maximize} & \sum_s U_s(x_s) \\ \text{subject to} & \sum_{s:l \in L(s)} x_s \leq \log(K\mathsf{SIR}_l(\mathbf{P})), \;\; \forall l, \\ & \mathbf{x}, \mathbf{P} \succeq 0 \end{array}$$

New Challenges:

- Global nonlinear (and nonconvex) dependency between rates and powers
- Need distributive algorithm to balance supply and demand

JOCP Algorithm

1. At each intermediate node, queuing delay λ is implicitly updated:

$$\lambda_l(t+1) = \left[\lambda_l(t) + \frac{\gamma}{c_l(t)}\left(\sum_{s:l\in L(s)} x_s(t) - c_l(t)\right)\right]^+$$

2. At each source, window size updated (and $x_s(t+1) = \frac{w_s(t+1)}{D_s(t)}$):

$$w_s(t+1) = \begin{cases} w_s(t) + \frac{1}{D_s(t)} & \text{if } \frac{w_s(t)}{d_s} - \frac{w_s(t)}{D_s(t)} < \alpha_s \\ w_s(t) - \frac{1}{D_s(t)} & \text{if } \frac{w_s(t)}{d_s} - \frac{w_s(t)}{D_s(t)} > \alpha_s \\ w_s(t) & \text{else.} \end{cases}$$

3. Each transmitter j passes message m_j to all other transmitters:

$$m_j(t) = \frac{\lambda_j(t)\mathsf{SIR}_j(t)}{P_j(t)G_{jj}}$$

4. Each transmitter updates its power:

$$P_l(t+1) = P_l(t) + \frac{\kappa\lambda_l(t)}{P_l(t)} - \kappa\sum_{j\neq l} G_{lj}m_j(t)$$

Example

A small example:

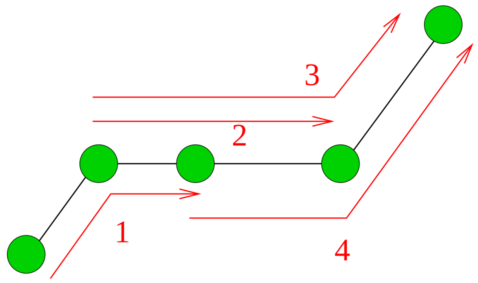

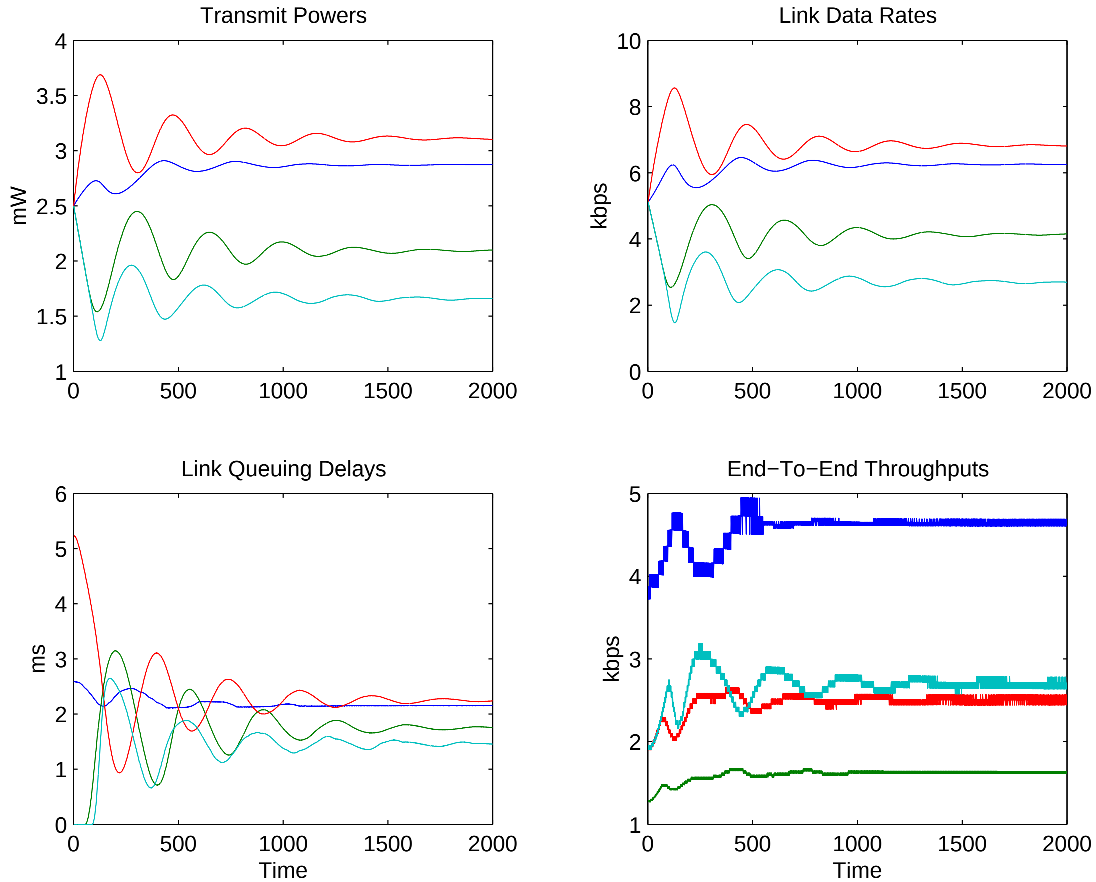

- 82% increase in end-to-end throughput per watt of power transmitted

- Most benefits from limited message passing

Layering Price

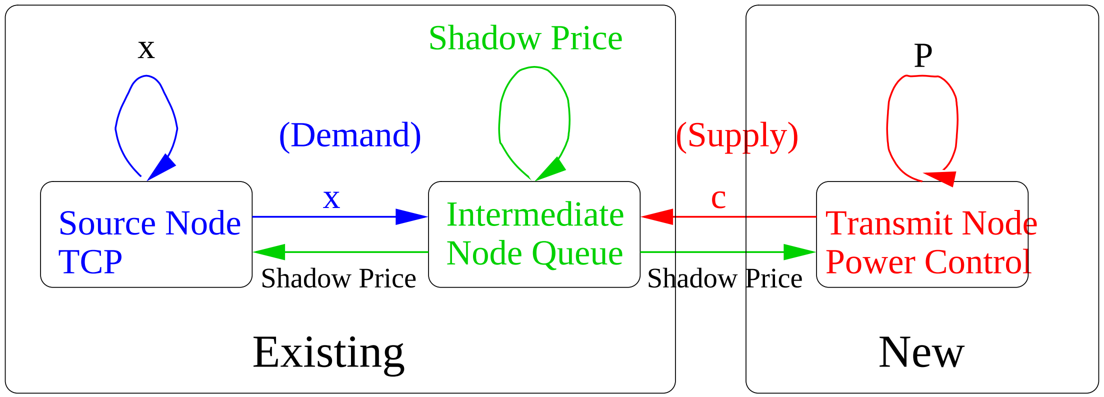

- Theorem: Nonlinearly coupled system converges to joint, global optimality
- Advantage: No need to change the existing TCP congestion control and queue management algorithms. Just utilize the values of queue length in designing power control algorithm in physical layer
- Congestion price is also layering price

Physical Layer Channel Coding: Formulation

Rate-reliability tradeoff. NUM in $(x_s, P_{e,s}, P_{e,l})$:

$$\begin{array}{ll} \text{maximize} & \sum_s U_s(x_s) + \sum_s V_s(P_{e,s}) \\ \text{subject to} & \sum_{s:l\in L(s)} x_s \leq c_l(P_{e,l}), \;\; \forall l, \\ & P_{e,s} = \sum_{l \in L(s)} P_{e,l}, \;\; \forall s \end{array}$$

- Source tradeoff: Higher rate, lower quality
- Link tradeoff: Fatter pipe, lower reliability

How to distributively find the network-wide, globally optimal tradeoff?

Physical Layer Channel Coding: Formulation

Rate-reliability tradeoff. NUM in $(x_s, P_{e,s}, P_{e,l})$:

$$\begin{array}{ll} \text{maximize} & \sum_s U_s(x_s) + \sum_s V_s(P_{e,s}) \\ \text{subject to} & \sum_{s:l \in L(s)} x_s \leq c_l(P_{e,l}), \ \ \forall l, \\ & P_{e,s} = \sum_{l \in L(s)} P_{e,l}, \ \ \forall s \end{array}$$

- Source tradeoff: Higher rate, lower quality
- Link tradeoff: Fatter pipe, lower reliability

How to distributively find the network-wide, globally optimal tradeoff?

Possible for large block length and reasonable decoding error probability

Physical Layer Channel Coding: Algorithm

- Source problem: maximize the sum of net utility on rate (with total congestion price) and net utility on reliability (with signal quality price)
- Source algorithm: local solution of source problem (2 variables), updates signal quality price

- Network problem: receive revenue from rate, pay price for unreliability
- Link algorithm: update link congestion price

Physical Layer Channel Coding: Algorithm

- Source problem: maximize the sum of net utility on rate (with total congestion price) and net utility on reliability (with signal quality price)
- Source algorithm: local solution of source problem (2 variables), updates signal quality price

- Network problem: receive revenue from rate, pay price for unreliability
- Link algorithm: update link congestion price

Next step: Recover information theoretic channel capacity as a special case of NUM: utility is rate and constraint is decoding error probability

Achievability proof as primal problems, converse proof as dual problems

Layering as Optimization Decomposition

How to layer? How not to layer? Separation theorem?

Rigorously quantify architectural principles and tradeoffs of layering

- Decompositions of a generalized NUM $\Leftrightarrow$ Layering schemes
- Variables coordinating the subproblems $\Leftrightarrow$ Interfaces among the layers
- Vertical and horizontal decomposition

Layering as Optimization Decomposition

How to layer? How not to layer? Separation theorem?

Rigorously quantify architectural principles and tradeoffs of layering

- Decompositions of a generalized NUM $\Leftrightarrow$ Layering schemes
- Variables coordinating the subproblems $\Leftrightarrow$ Interfaces among the layers
- Vertical and horizontal decomposition

Generalized Network Utility Maximization:

$$
\begin{array}{ll}
\text{maximize} & \sum_s U_s(x_s) + \sum_j V_j(w_j) \\
\text{subject to} & \mathbf{R}\mathbf{x} \preceq \mathbf{c}(\mathbf{w}, \mathbf{P}_e), \\
& \mathbf{x} \in \mathcal{C}_1(\mathbf{P}_e) \bigcap \mathcal{C}_2(\mathbf{F}), \\
& \mathbf{R} \in \mathcal{R}, \;\; \mathbf{F} \in \mathcal{F}, \;\; \mathbf{w} \in \mathcal{W}
\end{array}
$$

Optimization variables: $\mathbf{x}, \mathbf{w}, \mathbf{P}_e, \mathbf{R}, \mathbf{F}$

One Possibility

- Application layer. Utility functions U_i and V_j model application needs
- Transport layer. End-to-end throughput is source rate x_s for each end user s
- Network layer. Routing matrix can be designed by varying $\mathbf{R}$ within constraint set $\mathcal{R}$
- Link layer. Through scheduling, antenna beamforming, and spreading code assignment, contention matrix $\mathbf{F}$ can be designed within constraint set $\mathcal{F}$. Rates are then constrained by contention-free or contention-based access schemes described by constraint set $\mathcal{C}_2$.
- Physical layer. Adaptive resource allocations, *e.g.*, power control, adaptive modulation, coding with embedded diversity, leads to different logical link capacities $\mathbf{c}$ as functions of decoding error probabilities $\mathbf{P}_e$. Rate-reliability tradeoff forms constraint set $\mathcal{C}_1$

Are Alternative Decompositions Possible?

Consider a special case, with variables $(\mathbf{x}, \mathbf{y})$:

$$\begin{array}{ll} \text{maximize} & \sum_s U_s(x_s) \\ \text{subject to} & f_s(x_s, y_s) \leq 1, \\ & \sum_s g_s(y_s) \leq 0, \\ & x_s \in \mathcal{X}_s, \\ & y_s \in \mathcal{Y}_s \end{array}$$

- Wireless cellular network downlink power control
- Waterfilling algorithms for capacity computation
- Generalized coding problem for HOT systems

Are Alternative Decompositions Possible?

Consider a special case, with variables $(\mathbf{x}, \mathbf{y})$:

$$
\begin{array}{ll}
\text{maximize} & \sum_s U_s(x_s) \\
\text{subject to} & f_s(x_s, y_s) \leq 1, \\
& \sum_s g_s(y_s) \leq 0, \\
& x_s \in \mathcal{X}_s, \\
& y_s \in \mathcal{Y}_s
\end{array}
$$

- Wireless cellular network downlink power control
- Waterfilling algorithms for capacity computation
- Generalized coding problem for HOT systems

3 - 6 different decompositions depending on the applications

Lead to different tradeoffs between computation and communication, efficiency and complexity, different asymmetries in message passing

Remaining Challenges

Open Issue 1: Wireless networks: no decoding or multiuser decoding? MAC layer issues?

Open Issue 2: Internet routing protocols, especially BGP and wireless ad hoc routing?

Open Issue 3: Application layer modelling? Bursty and real-time traffic?

Open Issue 4: Functions that may not fit into optimization?

Open Issue 5: Non-convexity, integer constraints, nonzero duality gap for generalized network utility maximization?

Open Issue 6: Transient behaviors and time-scale mismatch?

Quantifying Network X-ities

Layering, and many networking principles, not just for performance metrics, *e.g.*, throughput, latency, distortion

Quantifying Network X-ities

Layering, and many networking principles, not just for performance metrics, *e.g.*, throughput, latency, distortion

It's for network robustness in terms of Network X-ities:

- Evolvability
- Scalability
- Deployability
- Adaptability
- Manageability

Quantifying Network X-ities

Layering, and many networking principles, not just for performance metrics, *e.g.*, throughput, latency, distortion

It's for network robustness in terms of Network X-ities:

- Evolvability
- Scalability
- Deployability
- Adaptability
- Manageability

Major challenges and opportunities:

- How to quantify these important and fuzzy notions?
- Optimize for the right objectives?

END Tool

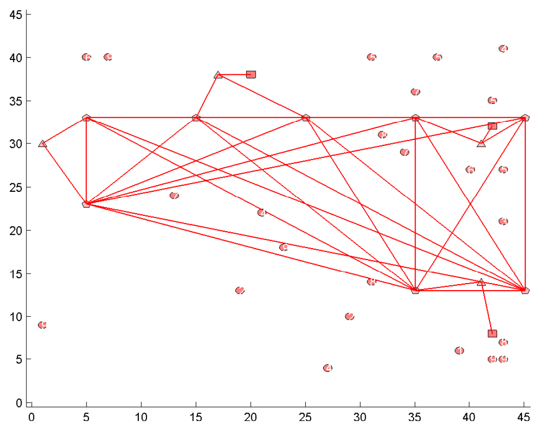
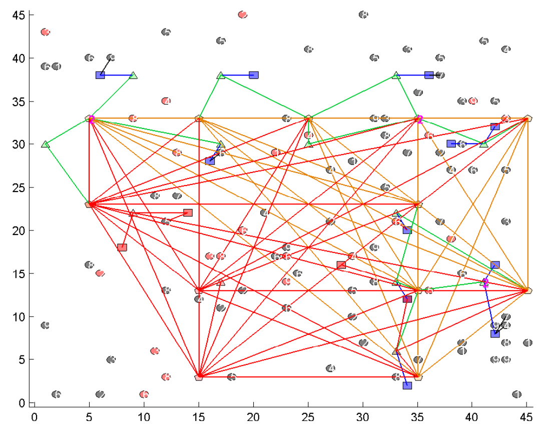
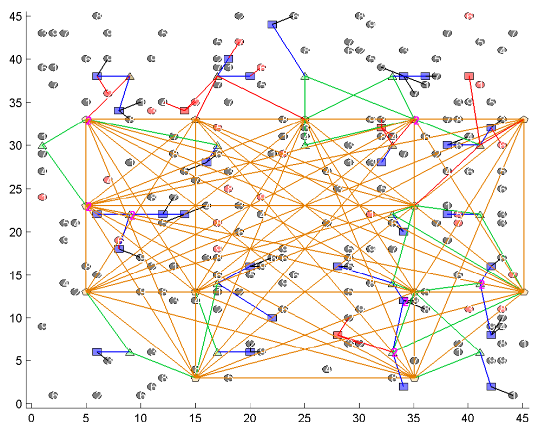
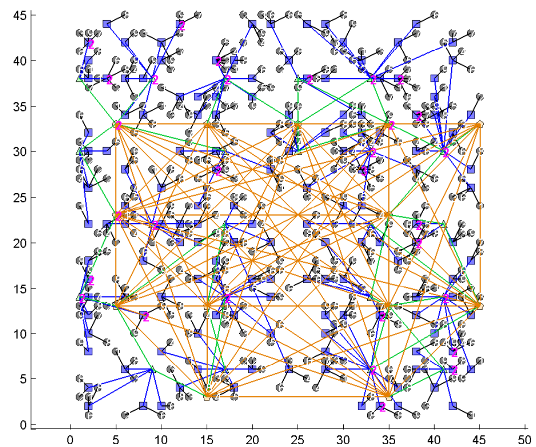

Summary

Tool: Network Utility Maximization

- An emerging, unifying framework for analysis and design of communication systems
- Substantial theoretical advances in recent years
- Significant practical motivations and applications

Goal: Layering As Optimization Decomposition

- Reverse engineering to sharpen understanding of layering interactions
- Forward engineering to design layering
- Many puzzles have recently been solved. Other challenges remain

Contacts

chiangm@princeton.edu
www.princeton.edu/∼chiangm